

Intel® 5100 Memory Controller Hub Chipset

Specification Update

May 2008

Revision 005US

Notice: The Intel® 5100 Memory Controller Hub Chipset may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are documented in this Specification Update.

Order Number: 318385-005US



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting [Intel's Web Site](#).

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Inside, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, IntelDX2, IntelDX4, IntelSX2, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, Viiv Inside, vPro Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2007–2008, Intel Corporation. All rights reserved.



Contents

Revision History	4
Introduction	5
Summary Tables of Current Product Issue Activity	7
General Product Information	12
Errata	13
Specification Changes	28
Specification Clarifications	29
Document-Only Changes	30



Revision History

Date	Revision	Description
May 2008	005	Added Erratum 44
April 2008	004	Added Erratum 43
March 2008	003	Clarified that Erratum 42 has no workaround
January 2008	002	Updated content for B0 silicon and DMA engine errata Added Errata 35 through 42 Updated workarounds for Errata 28 and 29 Updated General Product Information Added B0 stepping column to Table 6, "Errata"
September 2007	001	Initial release

Table 1. Revision Number Descriptions

Revision	Associated Life Cycle Milestone	Release Information
0.0	POP L3 Closure	Initial Documentation - Typically Internal Only
0.1–0.4	When Needed	Project Dependent - Typically Internal Only
0.5	Design Win Phase	First, Required Customer Release
0.6–0.7	When Needed	Project Dependent
0.7	Simulations Complete	Second, Recommended Customer Release
0.8–0.9	When Needed	Project Dependent
1.0	First Silicon Samples	Required Customer Release
1.1–1.4	When Needed	Project Dependent (Recommended)
1.5	Qualification Silicon Samples	Project Dependent
1.6–1.9	When Needed	Project Dependent
NDA - 2.0 Public - XXXXXX-001	First SKU Launch	Required Customer Release - Product Launch
2.1 and up	When Needed	Project Dependent

Note: Rows highlighted in gray are required revisions.



Introduction

Purpose/Scope/Audience

This document is an update to the specifications listed in the [Parent Documents/Related Documents](#) table that follows. This document is a compilation of [Errata](#), [Specification Changes](#), [Specification Clarifications](#), and [Document-Only Changes](#). It is intended for hardware and software system designers and manufacturers as well as developers of applications, operating systems, or tools.

Information types defined in [Conventions and Terminology](#) are consolidated into the Specification Update and are no longer published in other documents.

This document may also contain information that was not previously published.

Table 2. Parent Documents/Related Documents

Title	Number
Quad-Core and Dual-Core Intel® Xeon® Processor 5000 Sequence with Intel® 5100 Memory Controller Hub Chipset for Communications, Embedded, and Storage Applications – Platform Design Guide	Note 1
Intel® 5100 Memory Controller Hub Chipset Datasheet	http://www.intel.com/ (318378)
Intel® 5100 Memory Controller Hub Chipset (embedded) – External Design Specification (EDS) Addendum	Note 1
Intel® 5100 Memory Controller Hub Chipset for Communications, Embedded, and Storage Applications Thermal/Mechanical Design Guide	http://www.intel.com/ (318676)
RS - Intel® 5100 Memory Controller Hub Chipset BIOS Specification	Note 1

Notes:

1. Contact your Intel sales representative. Some documents may not be available at this time.


Conventions and Terminology

Note: [Errata](#) remain in the Specification Update throughout the product's lifecycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the Specification Update are archived and available upon request. [Specification Changes](#), [Specification Clarifications](#) and [Document-Only Changes](#) are removed from the Specification Update when the appropriate changes are made to the appropriate product specification or user documentation (datasheets, manuals, etc.).



Table 3. Conventions and Terminology

Term	Definition
Document-Only Changes	Document-Only Changes are changes to an Intel Parent Specification that result in changes only to an Intel customer document but no changes to a specification or to a parameter for an Intel product. An example of a document-only change is the correction of a typographical error.
Errata (plural) Erratum (singular)	Errata are design defects or errors. These may cause the Intel® 5100 Memory Controller Hub Chipset's behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.
Parent Specification	A parent specification is a top-level specification from which other documents can be derived, depending on the product or platform. Typically, a parent specification includes a product's pinout, architectural overview, device operation, hardware interface, or electrical specifications. Examples of parent specifications include the following: Datasheet, External Design Specification (EDS), Developer's Manual, Technical Product Specification. The derived documents may be used for purposes other than that for which the parent specification is used.
Specification Changes	Specification Changes are the result of adding, removing, or changing a feature, after which an Intel product subsequently operates differently than specified in an Intel Parent Specification , but typically the customer does not have to do anything to achieve proper device functionality as a result of Intel adding, removing, or changing a feature.
Specification Clarifications	Specification Clarifications are changes to a document that arise when an Intel Parent Specification must be reworded so that the specification is either more clear or not in conflict with another specification.



Summary Tables of Current Product Issue Activity

Table 6 through Table 9 indicate the [Errata](#), [Specification Changes](#), [Specification Clarifications](#), and [Document-Only Changes](#) that apply to the Intel® 5100 Memory Controller Hub Chipset product. Intel may fix some of the [Errata](#) in a future stepping of the component as noted in Table 4 or account for the other outstanding issues through [Specification Changes](#), [Specification Clarifications](#), or [Document-Only Changes](#). Table 6 through Table 9 use the codes listed in Table 4.

Table 4. Codes Used in Summary Tables

Code	Column	Definition
X	Stepping	Indicates either that, for the stepping/revision listed, <ul style="list-style-type: none"> an erratum eXists and is not yet fixed a specification change or specification clarification applies
No mark or blank	Stepping	Indicates either that, for the stepping/revision listed, <ul style="list-style-type: none"> an erratum is fixed a specification change or specification clarification does not apply
Plan Fix	Status	This erratum may be fixed in a future stepping/revision.
Fixed	Status	This erratum has been previously fixed.
No Fix	Status	There are no plans to fix this erratum.
A change bar to the left of a table row indicates an item that is either new or modified from the previous version of the Specification Update document.		

The errata inherited from the Intel® 5000 Series Chipset include Errata 1 through 27, 35, 36, 38 through 40, and 43. The cross-references between the Errata are shown in Table 5, “Intel® 5100 Memory Controller Hub Chipset to Intel® 5000 Series Chipset Errata Cross-reference”.

Table 5. Intel® 5100 Memory Controller Hub Chipset to Intel® 5000 Series Chipset Errata Cross-reference (Sheet 1 of 3)

Intel® 5100 Memory Controller Hub Chipset Erratum Number	Intel® 5000 Series Chipset Erratum Number ¹	Errata Title
1	1	PCI Express* auto link negotiation occasionally fails to correctly detect link width
2	2	MCH B1Err errors logged incorrectly
3	3	PCI Express* IBIST on x8/x16 port will not stop testing of entire port if an error is detected on any lane other than lowest four lanes of the port
4	6	PCI Express* receiver error logged upon putting link in disable state
5	7	PCI Hot Plug* ABP Bit Set While Attention Button is Pressed

Notes:

- Erratum, Compliance Issue, and Documentation Change reference numbers are based on the *Intel(R) 5000 Series Chipset Memory Controller Hub (MCH) Specification Update NDA, July 2007*.
- N/A = Not Available



Table 5. Intel® 5100 Memory Controller Hub Chipset to Intel® 5000 Series Chipset Errata Cross-reference (Sheet 2 of 3)

Intel® 5100 Memory Controller Hub Chipset Erratum Number	Intel® 5000 Series Chipset Erratum Number ¹	Errata Title
6	8	Leakage from 1.5 V VCC to 1.2 V VTT
7	9	Surprise Link Down Error Reported When PCI Express* Slot Power is Removed During Hot-plug Event
8	10	System hang with large number of transaction retries
9	11	INTL[7:2] registers are not implemented as read/write
10	13	SMI escalation via ERR[2:0]# pins may result in IERR#
11	15	CPU may record signal glitches when MCH is being reset
12	16	Coalesce mode cannot be used with maximum payload size of 256 B
13	22	MCH may log F2Err during shutdown special cycle initiated due to FSB timeout
14	25	Header of malformed TLPs on a x16 port not always logged
15	27	RID=CRID is sticky across warm reset
16	28	SLD could cause spurious completions
17	29	IBIST does not capture failed lanes properly
18	30	IBIST RX logic does not stop when IBISTR is reset
19	33	SLD on PCI Express* port during peer-to-peer posted requests can cause ESI to hang
20	34	PEXGCTRL.PME_TO_ACK may not be set when a turn off acknowledge TLP has been received from all ports
21	35	Masked completer abort status errors may be reported in the UNCERRSTS[0] register
22	36	SMBus TLOW:SEXT specification may be exceeded on SMBus 0 when the north bridge is clocked with a 266 MHz BUSCLK
23	37	First uncorrectable fatal bit of the root error status register may be incorrectly set when a second uncorrectable fatal error is received
24	38	Bit [3], INTxST, of the PEXSTS register may be cleared when bit [10], INTxDisable, of the PEXCMD register is set
25	N/A ²	Valid formatted TLP with invalid type may hang system
26	Compliance Issue #1	PCI Express* Compliance Test Failure: CBDMA Interrupt Line Register
27	Compliance Issue #2	PCI Express* Compliance Test Failure: MSI Address Register
35	41	PCI Express* transaction I/O ordering queue overflow
36	Document Change #2	PCI Express* device number assignment and header log
38	5	DMA engine channel completion address logged twice in FERR_CHANCOMP register

Notes:

1. Erratum, Compliance Issue, and Documentation Change reference numbers are based on the *Intel(R) 5000 Series Chipset Memory Controller Hub (MCH) Specification Update NDA, July 2007*.
2. N/A = Not Available



Table 5. Intel® 5100 Memory Controller Hub Chipset to Intel® 5000 Series Chipset Errata Cross-reference (Sheet 3 of 3)

Intel® 5100 Memory Controller Hub Chipset Erratum Number	Intel® 5000 Series Chipset Erratum Number ¹	Errata Title
39	26	Illegal addresses within 40-bit address space in channel completion address register do not generate cmp_addr_err
40	39	DMA engine's next channel error register not updated when subsequent errors detected
43	43	Read transactions may be delayed

Notes:

1. Erratum, Compliance Issue, and Documentation Change reference numbers are based on the *Intel(R) 5000 Series Chipset Memory Controller Hub (MCH) Specification Update NDA, July 2007*.
2. N/A = Not Available

Table 6. Errata (Sheet 1 of 2)

No.	Stepping/Revision		Status	Errata Title
	A0	B0		
1	X	X	No Fix	PCI Express* auto link negotiation occasionally fails to correctly detect link width
2	X	X	No Fix	MCH B1Err errors logged incorrectly
3	X	X	No Fix	PCI Express* IBIST on x8/x16 port will not stop testing of entire port if an error is detected on any lane other than lowest four lanes of the port
4	X	X	No Fix	PCI Express* receiver error logged upon putting link in disable state
5	X	X	No Fix	PCI Hot Plug* ABP Bit Set While Attention Button is Pressed
6	X	X	No Fix	Leakage from 1.5 V VCC to 1.2 V VTT
7	X	X	No Fix	Surprise Link Down Error Reported When PCI Express* Slot Power is Removed During Hot-plug Event
8	X	X	No Fix	System hang with large number of transaction retries
9	X	X	No Fix	INTL[7:2] registers are not implemented as read/write
10	X	X	No Fix	SMI escalation via ERR[2:0]# pins may result in IERR#
11	X	X	No Fix	CPU may record signal glitches when MCH is being reset
12	X	X	No Fix	Coalesce mode cannot be used with maximum payload size of 256 B
13	X	X	No Fix	MCH may log F2Err during shutdown special cycle initiated due to FSB timeout
14	X	X	No Fix	Header of malformed TLPs on a x16 port not always logged
15	X	X	No Fix	RID=CRID is sticky across warm reset
16	X	X	No Fix	SLD could cause spurious completions
17	X	X	No Fix	IBIST does not capture failed lanes properly
18	X	X	No Fix	IBIST RX logic does not stop when IBISTR is reset
19	X	X	No Fix	SLD on PCI Express* port during peer-to-peer posted requests can causes ESI to hang
20	X	X	No Fix	PEXGCTRL.PME_TO_ACK may not be set when a turn off acknowledge TLP has been received from all ports
21	X	X	No Fix	Masked completer abort status errors may be reported in the UNCERRSTS[0] register
22	X	X	No Fix	SMBus TLOW:SEXT specification may be exceeded on SMBus 0 when the north bridge is clocked with a 266 MHz BUSCLK



Table 6. Errata (Sheet 2 of 2)

No.	Stepping/ Revision		Status	Errata Title
	AO	BO		
23	X	X	No Fix	First uncorrectable fatal bit of the root error status register may be incorrectly set when a second uncorrectable fatal error is received
24	X	X	No Fix	Bit [3], INTxST, of the PEXSTS register may be cleared when bit [10], INTxDisable, of the PEXCMD register is set
25	X	X	No Fix	Valid formatted TLP with invalid type may hang system
26	X	X	No Fix	PCI Express* Compliance Test Failure: CBDMA Interrupt Line Register
27	X	X	No Fix	PCI Express* Compliance Test Failure: MSI Address Register
28	X		Fixed	Unpopulated ranks may cause performance degradation
29	X	X	No Fix	Flushing of writes in Coherency Engine (CE) for Asynchronous Self-refresh mode entry
30	X	X	No Fix	CKE does not go low during Asynchronous Self-refresh mode
31	X		Fixed	Write commands stuck in WR Queue preventing S3 entry in Asynchronous Self-refresh mode
32	X		Fixed	System failures while running memory sparing and patrol scrub
33	X		Fixed	Asynchronous Self-refresh relative to Reset deassertion
34	X	X	No Fix	Ports may be starved when both FSB buses are heavily loaded while PCI Express* agents are generating address conflicts
35	X	X	No Fix	PCI Express* transaction I/O ordering queue overflow
36	X	X	No Fix	PCI Express* device number assignment and header log
37		X	No Fix	Incorrect register default for proper DDR boundary scan board manufacturing testing
38	X	X	No Fix	DMA engine channel completion address logged twice in FERR_CHANCMR register
39	X	X	No Fix	Illegal addresses within 40-bit address space in channel completion address register do not generate cmp_addr_err
40	X	X	No Fix	DMA engine's next channel error register not updated when subsequent errors detected
41		X	No Fix	Incorrect CKE behavior during recovery from S3 suspend
42	X	X	No Fix	Machine check error due to a rare livelock condition when only two logical processors are enabled
43	X	X	No Fix	Read transactions may be delayed
44	X	X	No Fix	Disabling of Front Side Bus (FSB) parity impacts some data poisoning capabilities

Table 7. Specification Changes

No.	Stepping/Revision	Specification Changes
	AO/BO	
		None for this revision of the Specification Update.

**Table 8. Specification Clarifications**

No.	Stepping/ Revision	Specification Clarifications
	A0/B0	
1	X	PCI Hot Plug* Capability Validation

Table 9. Document-Only Changes

No.	Document Title	Rev.	Document-Only Changes
			None for this revision of the Specification Update.

General Product Information

The Intel® 5100 MCH Chipset can be identified by the following register contents:

Table 10. Identification Information

Part Number	Stepping	Vendor ID ¹	Device ID ²	Revision Number ³
Contact your Intel sales representative.	B0	8086h	65C0h	90h


Notes:

1. The Vendor ID corresponds to bits [15:0] of the VID - Vendor Identification Register located at Offset 00–01h in the PCI Function 0 configuration space.
2. The Device ID corresponds to bits [15:0] of the DID - Device Identification Register located at Offset 02–03h in the PCI Function 0 configuration space.
3. The Revision Number corresponds to bits [7:0] of the RID - Revision Identification Register located at Offset 08h in the PCI Function 0 configuration space.

The Intel® 5100 MCH Chipset stepping can be identified by the following component markings:

MCH	Stepping	Top Marking (GRP1LINE1)	QDF-Spec (GRP1LINE2) or MM#	Top Marking (GRP1LINE2)
Intel® 5100 MCH Chipset	B0	QG5100MCH	894163	Production

Figure 1. Top-Side Marking Example





Errata

The errata inherited from the Intel® 5000 Series Chipset include Errata 1 through 27, 35, 36, 38 through 40, and 43. The cross-references between the Errata are shown in Table 5, “Intel® 5100 Memory Controller Hub Chipset to Intel® 5000 Series Chipset Errata Cross-reference”.

1. PCI Express* auto link negotiation occasionally fails to correctly detect link width

Problem: The MCH occasionally fails to correctly detect the width of a hot-plugged PCI Express* card. This is restricted to the case where cards are “hot-plugged” with system power-on.

Implication: The PCI Express* auto link negotiation feature cannot be relied on to correctly detect the link width of a hot-plugged PCI Express* card during a hot add operation.

Workaround: Use the PEWIDTH[3:0] bits to set the width of the plugged in card. These bits are defined in the following table.

PEWIDTH[3:0]	Port0 (ESI)	Port2	Port3	Port4	Port5	Port6	Port7
0000	x4	x4	x4	x4	x4	x4	x4
0001	x4	x4	x4	x4	x4	x8	
0010	x4	x4	x4	x8		x4	x4
0011	x4	x4	x4	x8		x8	
0100	x4	x4	x4	x16			
others	Reserved						
1000	x4	x8		x4	x4	x4	x4
1001	x4	x8		x4	x4	x8	
1010	x4	x8		x8		x4	x4
1011	x4	x8		x8		x8	
1100	x4	x8		x16			
others	Reserved						
1111	x4	All port widths determined by link negotiation. ¹					

Notes:

1. Link negotiation configuration is not recommended due to an erratum.

Status: No Fix



2. MCH B1Err errors logged incorrectly

Problem: B1Err logging does not operate correctly. In the event that an internal data manager parity error occurs B1Err may not be logged. In addition, B1Err may get logged if poisoned data is passed to the internal data manager from the FSB and PCI Express* interfaces.

Implication: B1Err error reporting is unreliable.

Workaround: B1Err error logging should be disabled. The lack of B1Err error logging does not cause internal data manager errors to go undetected. An internal data manager parity error will be flagged at the destination interface.

Any parity error in the data manager will cause one of the following.

- A data parity error on the FSB if the destination is FSB
- A poisoned data pattern will be written to memory if the destination is system memory
- EP bit will be set on outbound PCI Express* packets if the destination is one of the PCI Express* ports.

Status: No Fix

3. PCI Express* IBIST on x8/x16 port will not stop testing of entire port if an error is detected on any lane other than lowest four lanes of the port


Problem: During PCI Express* IBIST “stop on error” testing, if an error is detected on a x4 port, the test will stop for the entire port, and the error is logged. Likewise, if an error is detected on the lower four lanes of a x8 port, the test will stop for the entire port, and the error is logged. If an error is detected on the upper four lanes of a x8 port, the error is logged, the test stops running on the upper four lanes, but the test will continue on the lower four lanes of the port.

Likewise, if an error is detected on the lowest four lanes of a x16 port, the test will stop for the entire port, and the error is logged. If an error is detected on any of the upper 12 lanes of a x16 port, the error is logged, the test stops running on the 4-lane group that the error occurred on, but the test will continue on the other 12 lanes of the port.

Implication: IBIST test will continue running in spite of errors if errors occur anywhere other than lowest four lanes of a x8/x16 port.

Workaround: If “stop on error” functionality is desired, changes will have to be made to the IBIST scripts to enable this functionality by polling the global status registers. If an error occurred in only the upper lanes, halt IBIST (if in loop continuous mode) as the IBIST engine for that set of four lanes has stopped (if stop on error is enabled). If an error occurred in the lowest four lanes, all of the IBIST engines will be forced to stop (exit from Loopback). Check the loop count error status register to determine what set of four lanes failed first. One thing to note is that the upper lanes will likely record an error in all lanes whereas the lower four may or may not depending on the nature of the real error. Whether the ports are in x4, x8 or x16 mode, the DIO0IBSTAT and DIO0IBERR registers need to be checked to determine which set of four lanes had an error and which set of four lanes did not.

Status: No Fix



4. PCI Express* receiver error logged upon putting link in disable state

Problem: When the link is told to go to Disable it is possible for the receive side to incorrectly log an error.

Implication: Receiver error incorrectly logged.

Workaround: To enter link disable, set the receiver error mask bit in EMASK_COR_PEX; then set the link disable bit. To exit link disable, clear the link disable bit; then clear the receiver error mask bit in EMASK_COR_PEX.

Status: No Fix

5. PCI Hot Plug* ABP Bit Set While Attention Button is Pressed

Problem: The Attention Button Pressed (ABP) bit (bit 0) in the PEXSLOTSTS register (Device 0, 2–7, Function 0, Offset 86h) is asserted while the hot-plug attention button is pressed. Software is supposed to clear this bit after the field has been read and processed but this clear operation could happen before the button is released, causing it to be set again.

Implication: Multiple attention button presses could be registered and processed during a single button press.

Workaround: Hot-plug drivers in Operating Systems capable of supporting native PCI Hot Plug* will not be supported. BIOS should use the _OSC method to maintain control of hot-plug events thereby only allowing ACPI hot-plug support. BIOS should implement a wait of 200 ms after the interrupt is received if using attention button. If using attention jumper, implement a wait of 2200 ms. Please contact your Intel representative to get the latest revision of the *RS - Intel® 5100 Memory Controller Hub Chipset BIOS Specification* for more details on the workaround.

Status: No Fix

6. Leakage from 1.5 V VCC to 1.2 V VTT

Problem: The *Quad-Core and Dual-Core Intel® Xeon® Processor 5000 Sequence with Intel® 5100 Memory Controller Hub Chipset for Communications, Embedded, and Storage Applications – Platform Design Guide* specifies that the MCH 1.5 V VCC power rail must ramp ahead of the MCH 1.2 V VTT power rail. During this power-up sequence, a leakage path exists within the MCH from the 1.5 V VCC pins to the 1.2 V VTT pins. This leakage path only exists while the MCH is powering up.

Implication: Higher than expected current seen on MCH 1.2 V VTT power rail while the 1.5 V VCC power rail is ramping. This leakage path does not impact the reliability or functionality of the MCH.

Workaround: None

Status: No Fix



7. Surprise Link Down Error Reported When PCI Express* Slot Power is Removed During Hot-plug Event

Problem: When power to a PCI Express* slot is removed during a hot remove event, a surprise link down error is incorrectly logged.

Implication: Surprise link down error is incorrectly logged.

Workaround: Hot-plug drivers in operating systems capable of supporting native PCI Hot Plug* will not be supported.

Perform the following.

1. When turning *off* the power to the slot:
 - a. Set bit [5] on the EMASK_UNCOR_PEX register.
 - b. Set bit [10] on the PEXSLOTCTRL register.
2. When turning *on* the power to the slot:
 - a. Clear bit [5] on the EMASK_UNCOR_PEX register.
 - b. Clear bit [10] on the PEXSLOTCTRL register.

BIOS should use the _OSC method to maintain control of hot-plug events, thereby only allowing ACPI hot-plug support. Please contact your Intel representative to get the latest revision of the *RS - Intel® 5100 Memory Controller Hub Chipset BIOS Specification* for more details on the workaround.

Status: No Fix

8. System hang with large number of transaction retries

Problem: Platforms can experience a system hang. This hang is characterized by a large number of transaction retries and repeated code fetches, as well as conflicting writes to the same address. No commercially available software has been observed to cause this condition in Intel's validation environment.

Implication: Possible system hang.

Workaround: The Memory Reference Code (MRC) includes a BIOS workaround that resolves the known instances of this erratum.

Status: No Fix

9. INTL[7:2] registers are not implemented as read/write

Problem: The INTL[7:2] registers are implemented as Read Only (RO) when they should be implemented as Read/Write (RW) as specified in the *PCI Local Bus Specification*, Rev. 2.3.

Implication: PCI compliance tests may report errors.

Workaround: None

Status: No Fix. See also Erratum 26.



10. SMI escalation via ERR[2:0]# pins may result in IERR#

Problem: Any condition routed to MCH ERR[2:0]# output pins for SMI# escalation that would result in data poisoning on FSB will result in a CPU IERR# assertion due to race condition. The race condition exists between poisoned data presented to the FSB by the MCH and the uncorrectable error escalated to SMI# via ERR[2:0]# pins. If the CPU observes an SMI# assertion before the poisoned data is presented to the FSB, the CPU asserts MCERR# followed by IERR#. The CPU expects the SMI# assertion to occur at least ten BCLKs after the associated MCERR# assertion. Conditions that result in data poisoning on the FSB are uncorrectable memory ECC errors or poisoned TLPs received from any PCI Express* port.

Implication: CPU may assert IERR# and cause a system hang before the error may be logged in SMM.


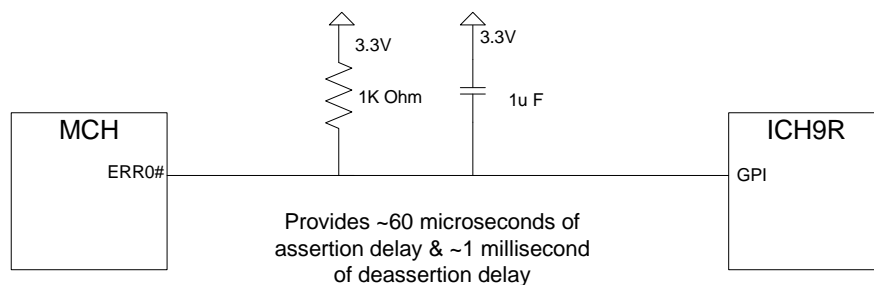
Workaround: Any one of the three workarounds that have been identified may be selected.

1. Log errors upon reboot

- System components (MCH, CPU, ICH9R) preserve errors across warm resets via sticky error registers
- Requires a mechanism to reset system upon IERR# (server management controller, depress system reset button, and so forth)

2. Delay SMI# assertion by inserting a 60–150 µs delay on the ERR0# pin

- Two recommended circuit solutions (other solutions may be possible):



- Route errors that result in FSB data poisoning to the ERR0# pin (IO4, M1)

3. Disable FSB data parity checking on CPU

- Prevents CPU from asserting MCERR# on poisoned data

Status: No Fix



11. CPU may record signal glitches when MCH is being reset

Problem: When the MCH is reset via RESETI# the CPU may record any one or more of the following errors: address strobe glitch (MSR IA32_MCI_STATUS bit [23]), data strobe glitch (MSR IA32_MCI_STATUS bit [22]), P/N data strobes out of sync (MSR IA32_MCI_STATUS bit [21]), PIC and FSB data parity (MSR IA32_MCI_STATUS bit [19]), RSP parity (MSR IA32_MCI_STATUS bit [18]), or FSB address parity (MSR IA32_MCI_STATUS bit [16]). This can happen when the MCH asserts CPURST# just after the MCH drives an FSB transaction. This may happen because RESETI# and CPURST# maintain an asynchronous relationship with each other.

Workaround: None

Status: No Fix

12. Coalesce mode cannot be used with maximum payload size of 256 B

Problem: When maximum payload size is set to 256 B (PEXDEVCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 74h, bits [7:5] set to '001'), and read completion coalescing is enabled (PEXCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 48h, bit [10] set to '1'), the system may hang. If MPS of 256 B is used, coalescing must be disabled. If coalescing is desired, MPS must be set to 128 B.

Implication: A system hang may occur.

Workaround: If an MPS of 256 B is desired, use the following settings.

1. Set maximum payload size to 256 B (PEXDEVCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 74h, bits [7:5] set to '001')
2. Disable coalescing (PEXCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 48h, bit [10] set to '0')

If coalescing is desired, use the following settings.

1. Enable coalescing (PEXCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 48h, bit [10] set to '1')
2. Set maximum payload size to 128 B (PEXDEVCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 74h, bits [7:5] set to '000')

In all cases, the following coalesce settings should be used.

1. Set COALESCE_MODE to 00 (PEXCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 48h, bits [25:24])
2. Use Max_rdcmp_lmt_EN default setting of 0 (PEXCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 48h, bit [12])
3. Use COALESCE_FORCE default setting of 0 (PEXCTRL[7:2,0] register, Device 0, 2–7, Function 0, Offset 48h, bit [11])

Note: Microsoft Server 2008* forces the maximum payload size to 256 B in the PEXDEVCTRL register; therefore, the workaround with coalescing disabled must be used for Microsoft Server 2008*.

Status: No Fix



13. MCH may log F2Err during shutdown special cycle initiated due to FSB timeout

Problem: The Intel® 5100 MCH Chipset will flag a FSB F2Err when a CPU issues a Shutdown Special Cycle due to FSB timeout with the deferred response enable (DEN#) not asserted. The CPU cycle is valid per FSB protocol.

Implication: An invalid FSB F2Err error is logged.

Workaround: There are two possible workarounds.

- BIOS may disable error logging for FSB F2Err by setting bit [1] of EMASK_FSB[1:0] registers (Device 16, Function 0, Offset 492h, 192h = 0n).
- Upon reboot error management software can check the CPU machine check registers for a FSB timeout status. If the FSB timeout status is asserted and the MCH is flagging F2Err, the F2Err assertion can be ignored.

Status: No Fix

14. Header of malformed TLPs on a x16 port not always logged

Problem: Quad-Core and Dual-Core Intel® Xeon® Processor 5000 Sequence with Intel® 5100 Memory Controller Hub Chipset for Communications, Embedded, and Storage Applications platforms may not log header information for an illegal length malformed TLP on a x16 port. This happens when the STP and END occur in the same symbol time - H H H E/IL - where H means Header DW, and E/IL is an END symbol with illegal length. Symbol time means the same clock cycle from the transaction layers point of view. For example, a x16 port will see four DWs of information in a single cycle.

Implication: HDRLOG registers may not contain the header of malformed TLP.

Workaround: None

Status: No Fix

15. RID=CRID is sticky across warm reset

Problem: The Revision Identification Register (RID) is keeping the value of Compatibility Revision ID (CRID) across a warm reset (RID=CRID) if RID is set to CRID prior to the warm reset. After a warm reset, RID should have the value of Stepping Revision ID (SRID) (RID=SRID).

Implication: When RID is set to show CRID's value and a warm reset occurs, RID will not be set to show SRID's value. RID will continue to show CRID's value. After a warm reset, it is expected that RID will show SRID's value.

Workaround: Prior to setting RID=CRID BIOS/ Software can store a copy of RID=SRID in a Sticky Scratch Pad Register, SPADS3[7:0] for example. After a warm reset BIOS/ Software can get the value of SRID from the Sticky Scratch Pad Register.

Status: No Fix

16. SLD could cause spurious completions

Problem: When a link goes down due to surprise link-down (SLD), the Intel® 5100 MCH Chipset may not drain all transaction layer packets (TLPs) before allowing the link to come back up.

Implication: In cases, where the Intel® 5100 MCH Chipset has not fully drained all downstream completions, a downstream device coming out of reset will see spurious completions.

Workaround: After an SLD, it is recommended that software clear the downstream device and root port of all errors.

Status: No Fix



17. IBIST does not capture failed lanes properly

Problem: The register that holds lane statistics is not sticky.

Implication: This means that only those lanes that detect an error last will be remembered. For example, if lanes 1 and 2 failed in IBIST cycle N and lanes 3 and 4 failed in cycle N+M, the error status register will only reflect lanes 3 and 4 as failed and will not show lanes 1 and 2.

Workaround: None

Status: No Fix

18. IBIST RX logic does not stop when IBISTR is reset

Problem: When IBIST is started in continuous mode and the IBISTR bit is cleared, the IBIST rx logic does not stop.

Implication: IBIST rx logic will continue comparing whatever data shows up at the input, even after the IBISTR bit is cleared.

Workaround: Clear the continuous mode bit before or along with the start bit.

Status: No Fix

19. SLD on PCI Express* port during peer-to-peer posted requests can causes ESI to hang

Problem: Under specific conditions a Surprise Link-down (SLD) could cause the ESI port to hang. For this to happen, a peer-to-peer write, with the source connected to the south bridge and destination connected to the north bridge, must be in progress. If an SLD occurs on the target link during certain phases of the MCH processing the request the ESI port may hang.

Implication: This bug causes the ESI port to hang which will hang the system.

Workaround: None

Status: No Fix

20. PEXGCTRL.PME_TO_ACK may not be set when a turn off acknowledge TLP has been received from all ports

Problem: When commanded to send PME Turn Off messages by setting bit [1], PME_TURN_OFF, of the PEXGCTRL register (Device 19, Function 0, Offset 17Ch) and a PME_TO_ACK TLP is received from a port before another port has sent its PME Turn Off message then bit [0], PME_TO_ACK, of the PEXGCTRL register (Device 19, Function 0, Offset 17Ch) may not be set when all the acknowledges have been received.

Implication: When a system is transitioning from S0 to S3, S4 or S5 a system hang may occur if the PCI Express* Global Control Registers Completion Timeout value has been set too low.

Workaround: Set bits [31:18], Completion Timeout, of the PEXGCTRL register (Device 19, Function 0, Offset 17Ch) to a value greater than 10 ms.

Status: No Fix

21. Masked completer abort status errors may be reported in the UNCERRSTS[0] register

Problem: When bit [15], IO7MSK, of the UNCERRMSK[0] register (Device 0, Function 0, Offset 108h) is set to mask Completer Abort status errors, uncorrectable Completer Abort status errors may still be logged in bit [15], IO7ERR, of the UNCERRSTS[0] register (Device 0, Function 0, Offset 104h) register.

Implication: A masked Completer Abort status error may be reported.

Workaround: None

Status: No Fix

**22. SMBus TLOW:SEXT specification may be exceeded on SMBus 0 when the north bridge is clocked with a 266 MHz BUSCLK**

Problem: The *System Management Bus (SMBus) Specification*, Version 2.0 TLOW:SEXT specification of 25 ms maximum, which applies only to Slave ports, may be exceeded on SMBus 0 by devices with a BUSCLK of 266 MHz.

Implication: A TLOW of ~31 ms may occur on SMBus 0. A Master is allowed to abort the transaction in progress to any slave that violates the TLOW:SEXT specification.

Workaround: None

Status: No Fix

23. First uncorrectable fatal bit of the root error status register may be incorrectly set when a second uncorrectable fatal error is received

Problem: Bit [4], FRST_UNCOR_FATAL, of the RPERRSTS register (Device 0, 2–3, 4–7, Function 0, Offset 130h) may be incorrectly set if an Uncorrectable Fatal Error is received and bit [2], ERR_FAT_NOFAT_RCVD, of the RPERRSTS register (Device 0, 2–3, 4–7, Function 0, Offset 130h) is set.

Implication: An Uncorrectable Fatal Error which was not the first error received may be incorrectly indicated as the type of error that was first received.

Workaround: None

Status: No Fix

24. Bit [3], INTxST, of the PEXSTS register may be cleared when bit [10], INTxDisable, of the PEXCMD register is set

Problem: When bit [10], INTxDisable, of the PEXCMD register (Device 0, 2–3, 4–8, Function 0, Offset 04h) is set to disable interrupts then bit [3], INTxST, of the PEXSTS register (Device 0, 2–3, 4–8, Function 0, Offset 06h) register may be cleared.

Implication: An interrupt status that is pending may be incorrectly cleared.

Workaround: None

Status: No Fix

25. Valid formatted TLP with invalid type may hang system

Problem: When a valid formatted TLP read is received, but it has an invalid type, it is incorrectly logged as a Master Abort Read transaction, updates the flow class status, and returns credit when it is supposed to ignore the transaction.

Implication: Returning the Credits creates a discrepancy between the MCH and PCI Express* partner. Eventually, the discrepancy in credits will cause a system hang. Receiving a valid formatted TLP with an invalid type is not expected to occur in a realistic system environment.

Workaround: None

Status: No Fix

**26. PCI Express* Compliance Test Failure: CBDMA Interrupt Line Register**

Problem: The PCI Express* compliance test TD_1_13 Interrupt Pin - Interrupt Line Register Test reports a RW failure for the interrupt Line Register bits [7:0] in Intel® 5100 MCH Chipsets. According to the *PCI Express* Base Specification*, Rev. 1.0a, Section 7.5.1.5, the Interrupt Line Register bits [7:0] should be RW. The INTL register is implemented as Read/Write Once (RWO) in Intel® 5100 MCH Chipsets. The register becomes RO after BIOS initializes it, violating the *PCI Express* Base Specification*, Rev. 1.0a. This is a minor violation because the DMA device does not have interrupt lines and INTL is not used by the Intel® 5100 MCH Chipset.

Implication: N/A

Workaround: N/A

Status: No Fix. Microsoft* Certification has decided not to make this a requirement for any OS WHQL. See also Erratum 9 in this document.

27. PCI Express* Compliance Test Failure: MSI Address Register

Problem: The PCI Express* compliance test TD_1_06 MSI Capability Structure Test reports a RO failure for the MSIAR bits [31:20] in Intel® 5100 MCH Chipsets. According to the *PCI Local Bus Specification*, Rev. 2.3, Section 6.8.1.4, Message Address for MSI, the Message Address Register bits [31:2] should be RW. The MSIAR in the Intel® 5100 MCH Chipset is a root port and is fixed to the Intel-specific IO_APIC range of FEEh to route the MSI ensuring proper functionality of the MSI architecture.

Implication: N/A

Workaround: N/A

Status: No Fix


28. Unpopulated ranks may cause performance degradation

Problem: Refresh cycles to unpopulated ranks may cause performance degradation depending upon data pattern and number of unpopulated ranks.

Implication: Performance degradation is possible up to approximately 5% on systems where less than all four ranks are populated. No degradation occurs if all four ranks are populated.

Workaround: The BIOS includes a workaround to resolve the known instances of this erratum.

Status: Fixed



29. Flushing of writes in Coherency Engine (CE) for Asynchronous Self-refresh mode entry

Problem: Write Queue not flushing completely when Asynchronous Self-refresh (Async Refresh) mode is enabled.

Implication: May not enter Self-refresh state when Async Refresh has been asserted.

Workaround: It is not trivial to determine whether an issued write transaction has been written to DRAM or not because such transactions are posted by the issuing agent. In order for application software to determine the last guaranteed successful write, the following algorithms must be applied depending on the origin of the data.

Note that the BIOS must set bit MCDEF3[12] to '1' in all cases.

FSB writes:

Each FSB agent must issue a locked read transaction (BLR) to any address (a config register for example) at periodic intervals (for example, after every 50 writes). Should an ASR occur, the locked read may or may not complete (return a value). In this case, the last write that was guaranteed to be successful is that immediately preceding the last locked read to complete.

I/O writes:

Each I/O agent must break the stream of writes into blocks of 96 or more writes, each block being followed by an inbound flush read to the address of the last (96th) write. The "blocks" are numbered in a sequence: 1, 2, 3, ..., N-1, N, N+1, ... In this case, block N+1 "guarantees" block N. In other words, if the flush read at the end of block N+1 completes successfully, all writes in block N (and preceding blocks) are guaranteed to have been executed. Note that completion of the flush read at end of block N does not guarantee any of the writes in block N (only those in N-1 and previous blocks).

DMA sequences:

Currently no mechanism exists for DMA sequences.

Writes to non-DRAM locations:

Currently no mechanism exists for posted writes to non-DRAM locations.

Status: No Fix

30. CKE does not go low during Asynchronous Self-refresh mode

Problem: CKE does not seem to be going LOW when the ASYNCRFSH pin is asserted during a heavy partial write to memory.

Implication: System may be prevented from entering S3 state.

Workaround: Set MCDEF3.ASRFRQ, register bit 11, when using the Asynchronous Self-refresh function.

Status: No Fix

31. Write commands stuck in WR Queue preventing S3 entry in Asynchronous Self-refresh mode

Problem: If a new write or partial request is sent to the MC during an Asynchronous Self-refresh entry to the S3 state, write commands may get stuck in WR Queue.

Implication: System may be prevented from entering S3 state.

Workaround: Under Investigation

Status: Fixed



32. System failures while running memory sparing and patrol scrub

Problem: The Intel® 5100 MCH Chipset may cause system failure when running spare copy with patrol scrub enabled.

Implication: May cause system hang or MCERR while running spare copy with patrol scrub enabled.

Workaround: None

Status: Fixed

33. Asynchronous Self-refresh relative to Reset deassertion

Problem: ASYNCRFSH pin assertion may not be detected at Reset deassertion, in other words, power loss just after reset deassertion.

Implication: The MCH may not transition into Asynchronous Self-refresh (Async Refresh) mode when a power down event occurs around Reset deassertion.

Workaround: Use CPLD logic to hold off ASYNCRFSH pin assertion until after Reset is deasserted, or hold ASYNCRFSH pin assertion until detected.

Status: Fixed

34. Ports may be starved when both FSB buses are heavily loaded while PCI Express* agents are generating address conflicts

Problem: If both the FSB ports are heavily loaded and the PCI Express* ports are trying to access the same address generating a lot of access conflicts due to FSB1 higher priority status, FSB0 and the PCI Express* ports may be starved.

Implication: FSB0 and PCI Express* ports may be prevented from accessing memory.

Workaround: Implement Patrol Scrub to prevent the condition.

Status: No Fix

35. PCI Express* transaction I/O ordering queue overflow

Problem: Under some corner case scenarios when the system is stressed with heavy I/O traffic, a hang condition could occur as the transactions to/from the PCI Express* port are blocked due to an I/O order queue overflow.

Implication: The overflow condition will eventually cause the system to deadlock because many outstanding transactions are unable to make forward progress. The system level implication is a system hang with IERR or MCERR or PCI Express* fatal error asserting NMI, if enabled.

Workaround: Please refer to the latest *RS - Intel® 5100 Memory Controller Hub Chipset BIOS Specification* release.

Status: No Fix

36. PCI Express* device number assignment and header log

Problem: The requestor ID field in the TLP of ports 4, 5, 6, 7 (devices 4–7) are incorrectly mapped as device 12, 13, 14, 15 (decimal), respectively, for non peer-to-peer and non-DMA sent outbound requests issued by the Intel® 5100 MCH Chipset.

Implication: The PCI Express* HDRLOG registers, which record TLP header information during I/O errors (viz. IO1Err, IO4Err, IO5Err, IO7Err, IO8Err, IO9Err, IO10Err) for completion TLPs returning to the MCH, will have the translated device numbers and not the original device number assignment.

Workaround: None

Status: No Fix



37. Incorrect register default for proper DDR boundary scan board manufacturing testing

Problem: Board manufacturing boundary scan testing of the DDR2 I/O by register default is disabled.

Implication: Prior to BIOS execution, if board manufacturing boundary scan testing of the DDR2 I/O interface is required, it will not function.

Workaround: Write the DRAMISCTL[1:0] register, VOXSTART bit to the correct value using the San_Clemente_MCH_Chipset_B0_DDR_boundary_scan_workaround_Rev1_0.svf Serial Vector File (SVF) under the conditions below, as also described in the San_Clemente_MCH_Chipset_B0_BSDL_Rev1_0_ReleaseNote.txt text file.

1. CORE_CLKP/N must be available to the MCH.
2. RESETI# must be de-asserted.
3. POWERGOOD must be asserted and remain asserted for the duration of the boundary scan test. If POWERGOOD is de-asserted, rewrite the register using the SVF file to re-enable DDR boundary scan.

All necessary information and files are included in the *Intel® 5100 Memory Controller Hub Chipset B0 Stepping (embedded) – Boundary Scan Description Language (BSDL) File*.

Status: No Fix

38. DMA engine channel completion address logged twice in FERR_CHANCMP register

Problem: A 32-bit channel completion address is logged in both the high order 32 bits and the low order 32 bits of the 64-bit FERR_CHANCMP register in the DMA engine.

Implication: The upper 32 bits of the 64-bit completion address are not logged.

Workaround: The correct upper 32 bits of the FERR_CHANCMP register can be obtained by getting the channel number logged from the FERR_CHANSTS.FERR_DMA Channel Number field (Device 8, Function 0, Offset 84h, bits [4:3]) and reading the CHANCMP (Device 8, Function 1, Offset 218h, 198h, 118h, 98h) register for that particular channel. This should cover all usage models, unless software reprogrammed the CHANCMP register after the error occurred and before this reading occurred.

Status: No Fix

39. Illegal addresses within 40-bit address space in channel completion address register do not generate cmp_addr_err

Problem: The Intel® 5100 MCH Chipset DMA engine will not flag a cmp_addr_err (DMA12) when illegal addresses within the 40-bit address space are programmed in the channel completion address register.

Implication: The cmp_addr_err (DMA12) bit will not get set when illegal addresses within the 40-bit address space are programmed in the channel completion address register. The error is flagged correctly when an address greater than 40 bits is programmed in the channel completion address register. However, in all cases where an illegal address is programmed, the illegal address will be caught, and the transaction will be master aborted.

Workaround: None

Status: No Fix

40. DMA engine's next channel error register not updated when subsequent errors detected

Problem: The NERR_CHANERR register (Device 8, Function 0, Offset BCh) is not updated when additional errors are detected; only the second error is logged.

Implication: The latest error status is not captured.

Workaround: None

Status: No Fix

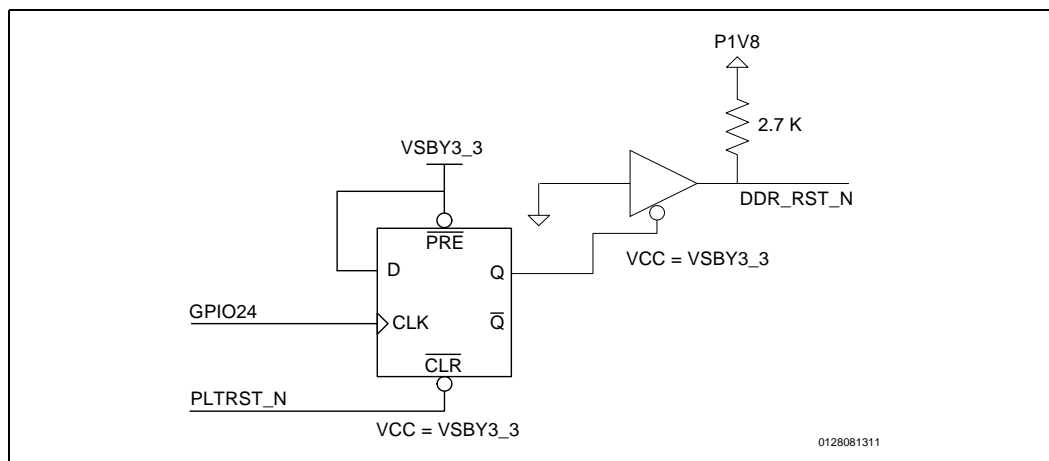
41. Incorrect CKE behavior during recovery from S3 suspend

Problem: During the initial stage of exit from System Sleep State "S3" (suspend to RAM), the CKE signal is driven to an indeterminate state (High, Low, or tristate).

Implication: Memory contents may not be preserved through S3 suspend, making S3 mode usage undependable. This includes implementations of the Asynchronous Request for Self-refresh feature.

Workaround: The workaround requires both a board-level circuit and a BIOS change. The workaround holds the DIMMs in reset using a flipflop clocked by a GPIO from the ICH9R under control of the BIOS. An output controlled buffer is also required to transition the 3.3 V output of the flipflop to the 1.8 V reset input of the DIMM (see diagram below). The possible GPIOs are GPIO[28,27,24,12]. The GPIO implemented in the circuit is incorporated into the BIOS as a build option by an Intel BIOS Vendor (IBV).

In the example below, GPIO24 is used to clock the flipflop. The GPIO output only controls the DIMM reset during S3 operation; any transitions on the GPIO after S3 exit have no effect on the circuit or the DIMMs Reset signal. The workaround potentially allows the GPIO to be used for other purposes. For more information, refer to the *Intel® 5100 Memory Controller Hub Chipset B0 Stepping (embedded) – S3 Erratum Workaround White Paper*, document number 379825.



Note: PLTRST_N can be either of the Platform Reset signals PLTRST1_N or PLTRST_N, DDR_RST_N is the DIMM reset signal, P1V8 is the DIMM supply, and VSBY3_3 is the standby 3.3 V supply (optionally VSBY5 standby) as referenced in the Tionesta and Williamsport CRB schematics.

Status: No Fix



42. Machine check error due to a rare livelock condition when only two logical processors are enabled

Problem: The Intel® 5100 MCH Chipset can reach a livelock situation when running two logical processor cores in either single socket (two cores enabled in one socket) or dual socket (single core enabled per socket) configurations. This happens if one logical core issues a zero length write invalidate and the other core issues a read transaction to the same cache line address.

Implication: This situation creates a livelock condition resulting in a machine check error (MCA) from the processor which, depending on BIOS handling of machine check errors, may result in a system hang or an automatic reboot of the system.

At this time, Intel has not been able to reproduce this erratum with existing commercial software. A primary reason is that multitasking operating systems use timers to periodically interrupt application threads for task switching purposes. A task switch would break the sequence and allow escape from a livelock condition.

Workaround: None

Status: No Fix

43. Read transactions may be delayed

Problem: Under a certain sequence of read and write transactions issued from processors or bus mastering I/O devices, the read transaction may be delayed.

Implication: A read transaction may be delayed. Intel has not observed this behavior with any commercially available software.

Workaround: None

Status: No Fix

44. Disabling of Front Side Bus (FSB) parity impacts some data poisoning capabilities

Problem: If the Intel® 5100 MCH Chipset is configured to disable Front Side Bus (FSB) parity, certain data transactions will not support poisoning of PCI Express* packets.

Note: Chipset FSB parity disabling is only required for specific CPUs lacking FSB parity signaling. Platforms using the Intel® Xeon® Processor 5000 Sequence do utilize FSB parity and are, therefore, unaffected.

Implication: Although certain peer-to-peer transactions can result in PCI Express* endpoints receiving non-poisoned packets, the CPU can still detect the bad data in the I/O interfaces or memory subsystem by way of the MCH error signals ERR[2:0]. If necessary, it will be up to error-handling software to appropriately purge or restore system data. Please contact your Intel representative to get the latest revision of the *RS - Intel® 5100 Memory Controller Hub Chipset BIOS Specification* for more details on disabling FSB parity and error handling.

Workaround: None

Status: No Fix



Specification Changes

None for this revision of this Specification Update.



Specification Clarifications

1. PCI Hot Plug* Capability Validation

The Intel® 5100 Memory Controller Hub Chipset currently provides limited hardware and documentation support for the PCI Hot Plug* capability. However, Intel validation efforts of the PCI Hot Plug* capability and public collateral are not expected to complete until 2H'08.



Document-Only Changes

None for this revision of this Specification Update.